Tagwebs, Flickr, and the Human Brain

An essay by Jakob Lodwick on February 1st, 2005

Tagging, as seen on Flickr and other sites, allows you to organize things in a way that makes sense to your brain. You may also notice that tagging photos on Flickr is the first time that organizing something has *made perfect sense.* According to Scientific American, in 1966 Ben-Ami Lipetz concluded that:

> **...breakthroughs in information retrieval would come when researchers gained a deeper understanding of how humans process information and then endowed machines with analogous capabilities.**

Well, Ben was right, as you'll soon see for yourself. By looking at how we tag photos on Flickr, we can understand how humans process information. Once we understand that, we can understand how to model it with computers, thereby creating better information retrieval systems.

What Ben was unable to predict all those years ago was that we will not only develop better information retreival systems, but also model our own brains on the lowest levels, and eventually create artificial intelligence.

An introduction to tagwebs:

What if we could tag not just photos, but also other tags? We could start to build a tagweb. When a tagweb is created from your tags, that tagweb works perfectly within the realm of what makes sense to you. The reason nobody came up with this before Flickr was because we didn't have Flickr as a visible reference point. You can't just *imagine* something out of the blue without first thinking about related things.

- A tagweb is a network of ideas that can be graphed on a computer or on paper.
- A tagweb is not built deliberately; it emerges naturally as users contribute to a system.
- A tagweb requires no external information.
- A tagweb emulates its creator's brain.
- As a tagweb grows, it becomes more useful.
- A tagweb does not "understand" its contents; rather, it understands the *relationships* of its contents.

Let's make a tagweb! Imagine that you had a photo of an adorable baby wearing giant nerd glasses. What would you tag this with? Well, when I look at that photo, a few ideas jump to mind: Victoria, glasses, and cute. It makes sense to me, so that's what I'll tag it with:

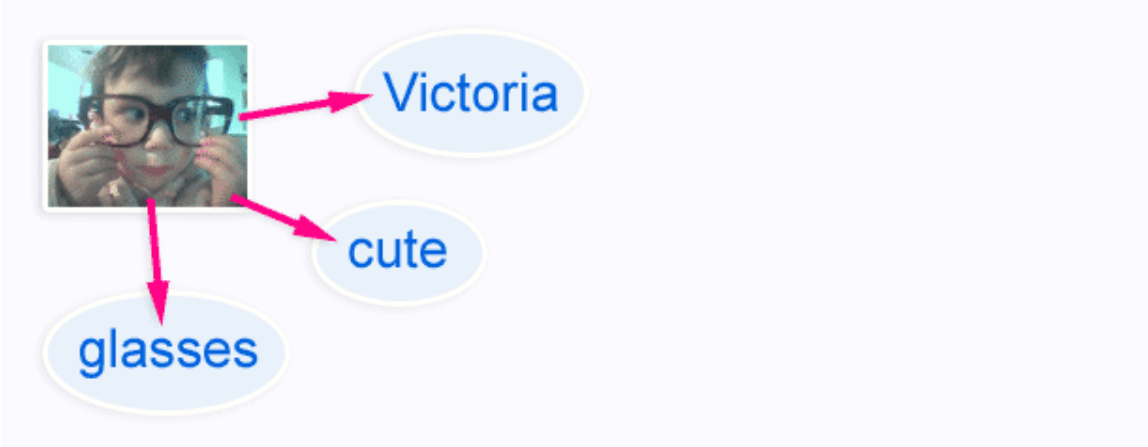

Flickr, you are truly awesome. Now, when somebody wants to see glasses or cute things or things called "Victoria", he or she might just find my photo. But what if someone had wanted to find photos of all sorts of females, or of people in general? They would have come up empty-handed, because the words "female" and "people" aren't directly tagged onto this image.

So, from now on, the tagweb you see is one that Flickr cannot create, but that I would create if I had the ability to tag my tags. If I could tag my tags, then I would tag the word "Victoria" with "female".

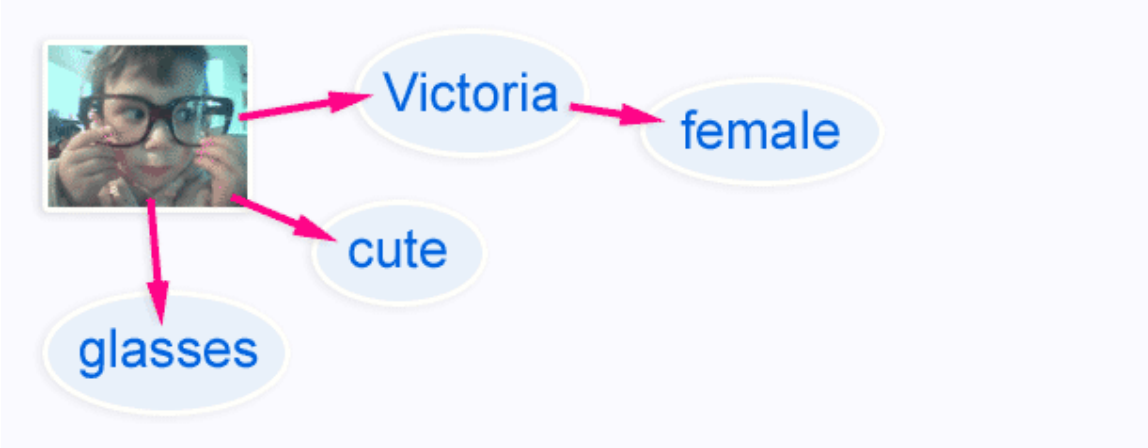

Now I'll tag "female" with "person", because if I was thinking about the idea of "female", one thing that comes to mind is "person". I'll also tag "Victoria" with "cute", even though I already tagged the photo with "cute", since I think both Victoria and the photo itself are cute.

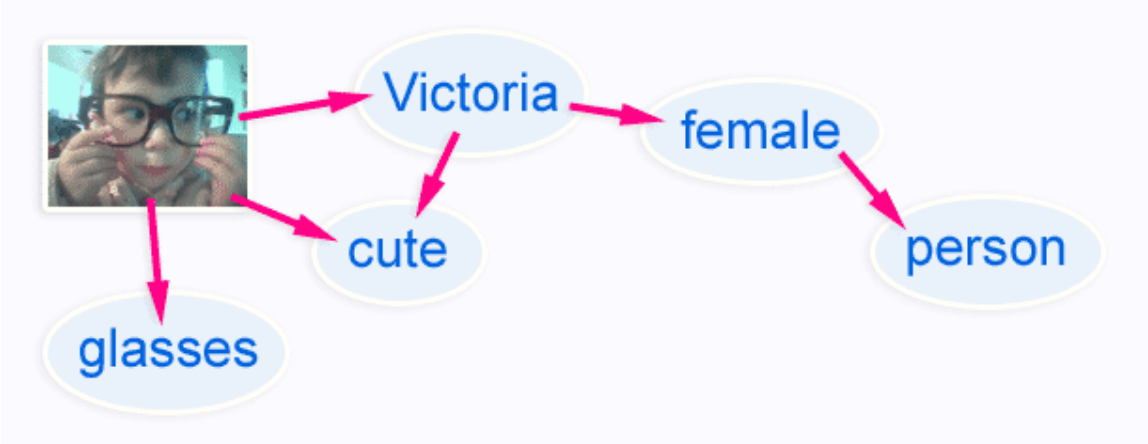

We now have an expanded tagweb, one that can only exist in theory because I cannot tag my tags on Flickr. Let's continue by adding a photo of myself holding a potted plant. As always, I will tag this

photo with terms that make sense to me - terms that I'm almost *compelled* to enter, even though

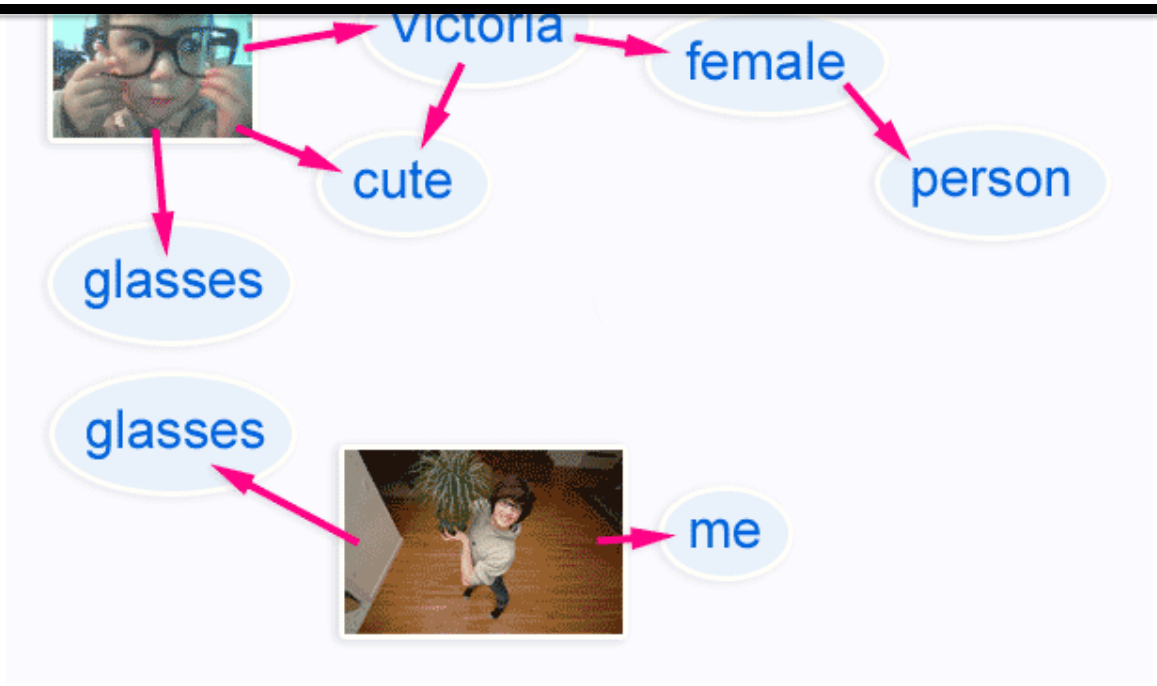

And, of course, I'll tag those new tags. And yes, like everyone in the world, I consider myself cute.

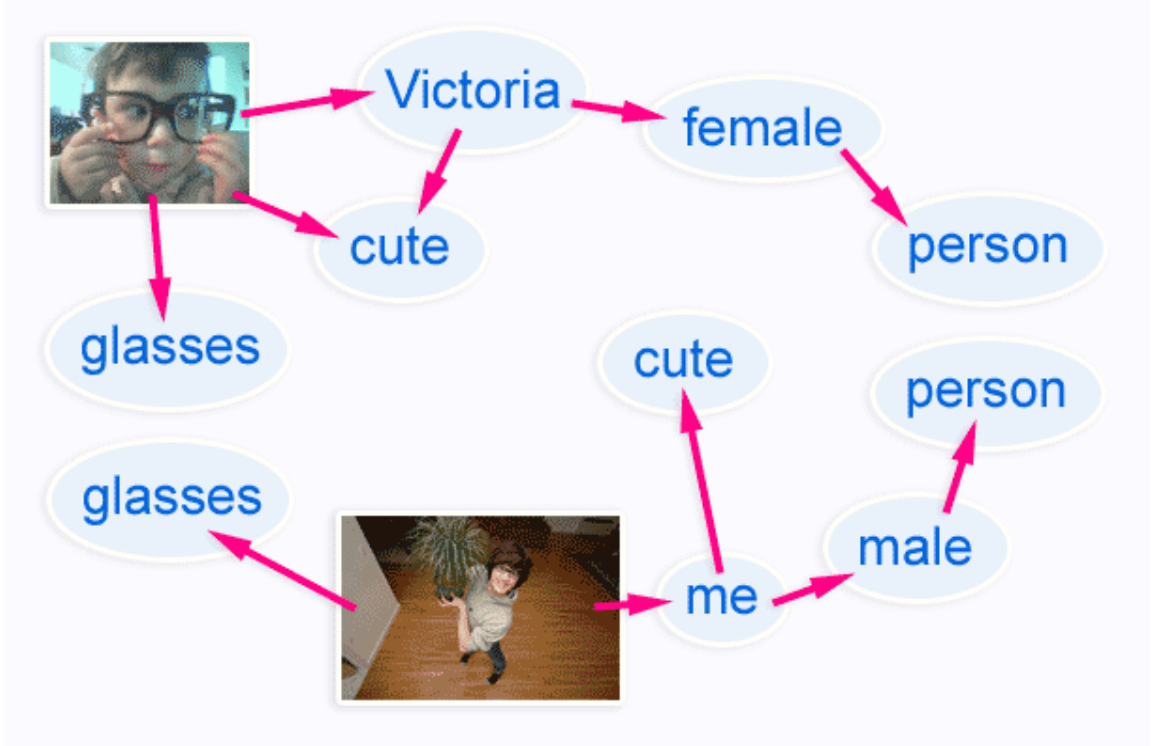

Now we have these two separate little tagwebs, each stemming from one photo. Why don't we consolidate this graph, and instead of having two "glasses" and "person" tags, just draw one of each?

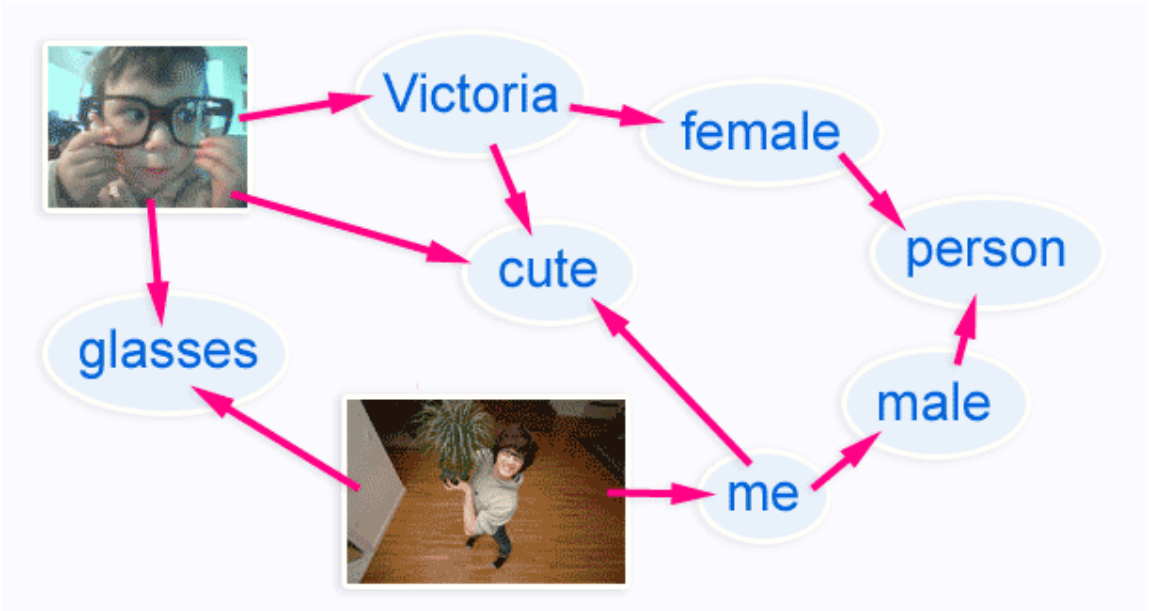

How relatively concise! Let's look at some interesting characteristics of the network as it now stands. What do we notice?

(this is a QuickTime movie called tagwebs-01.mov | the sound is low, turn up your speakers)

1. With regard to "person", "female" and "male" are connected in the same way. In other words, "female" and "male" both have the same relationship with "person".
2. "Me" has the same relationship with "male" as "Victoria" has with "female". As a simple analogy, Me:male::Victoria:female. The fact that the system can convey this information *without any idea of the meaning behind the words* is what excites me.
3. Me:female::Victoria:male. This analogy may sound stupid at first, but I think it's pretty interesting, and especially great that *the system itself* has this information stored inherently. A computer could tell you this analogy! It may not know how to describe that analogy in words, but the tagweb still knows it.
4. The word "person" has something to do with the word "glasses", but when you look at the relationship between those words, it always involves the idea of a person who wears them.

But what's the proof that a computer could tell me this? How could a system possibly "know about an analogy"? Just watch the video!

(this is a QuickTime movie called tagwebs-02.mov | the sound is low, turn up your speakers)

Let's make a graph comparing several of the relationships that make clear sense to us as English-speakers. Let's also include the pattern of each relationship. The pattern is determined by starting at the first idea and tracing a path to the second, and recording every arrow you pass through to get there. An arrow that moves with your path is marked by a **+**, and an arrow that moves against your path is marked with a **-**.

Words to compare	*Path taken*	*Pattern*	*English approximation*
Victoria :: female	Victoria → female	**+**	"Victoria is a female."
female :: person	female → person	**+**	"A female is a person."
male :: person	male → person	**+**	"A male is a person."
me :: male	me → male	**+**	"I am a male."
female :: male	female → person ← male	**+ -**	"Females and males are both examples of people, just as Jakob and Victoria are both examples of cuteness."
Victoria :: me	Victoria → cute ← me	**+ -**	
Victoria :: person	Victoria → female → person	**+ +**	"Victoria is a type

			of person in the same way that I am a type of person."
me :: person	me → male → person	+ +	
Victoria :: male	Victoria → female → person ← male	+ + -	"Victoria's gender is not male, just as my gender is not female."
me :: female	me → male →person ← female	+ + -	

By creating a pattern from the paths between two ideas, we can quantify the relative relationship of any two items in the network. We can define the relationship between "Victoria" and "male" as **+ + -**, which a computer could tell us is the same as the **+ + -** between "me" and "female".

The symbols are arbitrary; it doesn't matter if we're using **+** and **-** or **Bert** and **Ernie**. Even the direction of the arrows is arbitrary, as long as you're consistent. Try making a tagweb yourself, and seeing how different relationships compare. You'll notice that the larger you build your tagweb, the "smarter" it is and the more capable it is of revealing interesting analogies. Also, you'll notice that even if you try to "break" the tagweb by feeding it some faulty information, you'll still produce a system that works in many ways. This explains why a good architect can be a terrible poet, or why I'm able to program PHP but unable to comprehend football.

Q: What are the implications of these findings?

A: We now have an analogy for how our brains work on the lowest level. The analogy is simple: A neuron in your brain is a lot like a tag in a tagweb.

You understand something when you make an analogy of it. Imagine trying to explain a modern computer to someone from the 1400s. It would take a long time. But to explain a modern computer to a man from 1950 who already knew about television and typewriters, he'd catch on much quicker.

So, likewise, now that we have an understanding of "what neurons are like" and, by extension, a decent model of how our brains work, we can start doing things in ways that make much more sense. Our input and our output will be of higher quality. Conversation will improve, language barriers will fall, artificial intelligence will begin to emerge, and, hopefully, people will be more honest about what they want and who they want to have sex with.

Q: Who came up with these ideas?

I recently had a "Eureka!" moment laying in bed, drifting through my thoughts, when the concept of tagging a tag popped into my head. This was at 2:00am. For the next three hours, I frantically wrote out pieces of paper that explained my early thoughts on these ideas. These pieces of paper still make just as much sense to me now as they did then.

I can say with some confidence that I'm not the first person to understand these theories. I am fairly certain that Google is well aware of these ideas, which are closely related to their PageRank technology, and which could explain why they're able to do so many incredible things at such a feverish pace. (Disclosure: I own stock in GOOG).

I also suspect that the minds of Michael Gondry, Malcolm Gladwell, the late Jim Henson, and Warren Buffett all relate to the world in a way that approaches or embraces these ideas.

Q: Do you really think people will want to enter all this information?

Yes. Look how people are using Flickr already. When a user is looking at one thing, he is willing to tag it with a few words. This isn't asking too much, nor is it "metadadaism". Tagging something is the easiest way to enter metadata on something. Giving something a tag is a lot more important than giving it a title or caption. And giving a tag to something that was in a photo you took - something that is meaningful to you - is a fairly natural act. I'm sure some other people experience **The Flickr Effect**.

Also, let's say that only 5% of Flickr users are tagging photos. Five percent is a low amount, but look at how much that 5% has improved the system. Look at how easy it is to find dozens of photos of mohawks! Look how easily someone can get a decent idea of what I look like and who I hang out with!

Even if only one tenth of one percent of Flickr users start tagging their tags, it will still increase the value of the site noticeably. Even if only *I* start tagging *my* tags, we'll see an increased value.

Q: How does this relate to things like ConceptNet, WordNet, and Porter Stemming?

Because those things were invented before Flickr and del.icio.us, they had no viable metaphor for the way our brains work, and were therefore limited in scope. Tagwebs can explain the brain in a very thorough way, rather than explaining just words themselves.

Also, they put too much emphasis on the words themselves, when really there needs to be **no emphasis on words**. If you build a system that puts importance on *the structure of the word*, you're creating a system based around your own language. You must create a system based around your own *thoughts* if you want it to simulate your own thoughts. Keep in mind that the tagweb we created earlier could equally reflect the same relationships if the words were substituted with arbitrary symbols, the way a computer would process them:

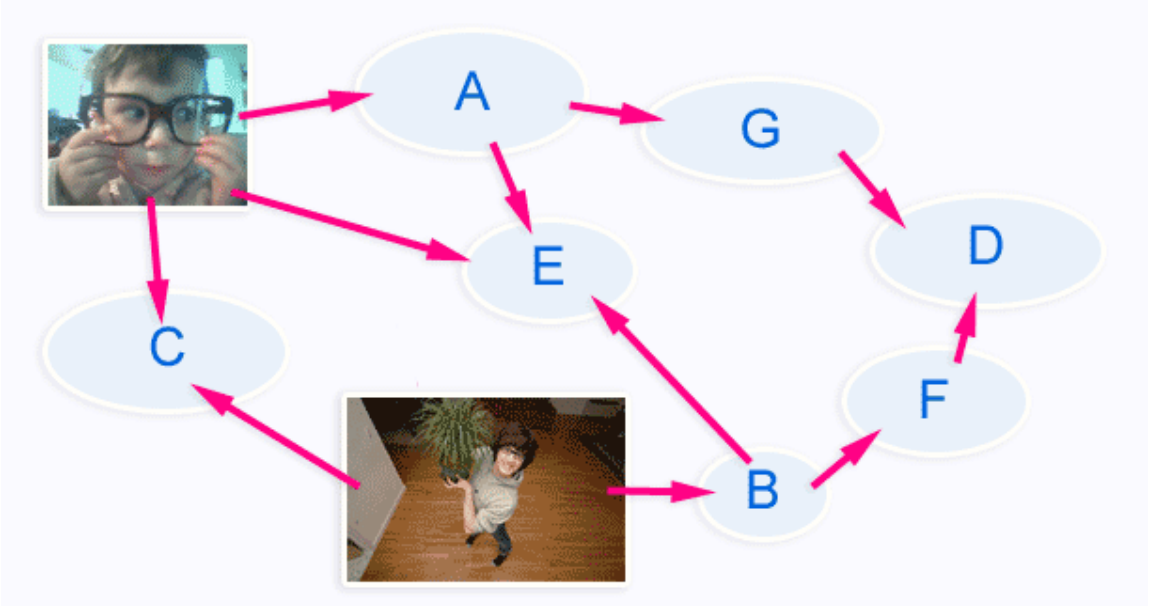

Conclusion (for now)...

Even though it's only been in my head for the past few days, the theory of tagwebs and the ideas that comprise it have fundamentally changed my brain. My abilities to solve problems, to make analogies, to make jokes, to explain tagwebs, to understand people, to *think* -- all of these things have been improved greatly. I now understand how my brain works, and I can act in ways that embraces that knowledge. I believe that we are on the verge of a major evolutionary step in the way we use and understand our brains, and that we will begin to see the effects of this change immediately.

> **"We are the universe trying to understand itself" - Carl Sagan**

Additionally, if you made it all the way through this article and would like to talk to me, then I would like to talk to you. My email address is listed on my home page. Or, you can discuss this idea in the FlickrIdeas forum.